

# End-to-End Cross-Modality Retrieval with CCA Projections and Pairwise Ranking Loss

Matthias Dorfer<sup>1</sup>, Jan Schlüter<sup>2</sup>, Andreu Vall<sup>1</sup>, Filip Korzeniowski<sup>1</sup>, and Gerhard Widmer<sup>1,2</sup>

<sup>1</sup>Department of Computational Perception, Johannes Kepler University Linz

<sup>2</sup>The Austrian Research Institute for Artificial Intelligence

April 17, 2018

## Abstract

Cross-modality retrieval encompasses retrieval tasks where the fetched items are of a different type than the search query, e.g., retrieving pictures relevant to a given text query. The state-of-the-art approach to cross-modality retrieval relies on learning a joint embedding space of the two modalities, where items from either modality are retrieved using nearest-neighbor search. In this work, we introduce a neural network layer based on Canonical Correlation Analysis (CCA) that learns better embedding spaces by analytically computing projections that maximize correlation. In contrast to previous approaches, the CCA Layer (CCAL) allows us to combine existing objectives for embedding space learning, such as pairwise ranking losses, with the optimal projections of CCA. We show the effectiveness of our approach for cross-modality retrieval on three different scenarios (text-to-image, audio-sheet-music and zero-shot retrieval), surpassing both Deep CCA and a multi-view network using freely learned projections optimized by a pairwise ranking loss, especially when little training data is available (the code for all three methods is released at: [https://github.com/CPJKU/cca\\_layer](https://github.com/CPJKU/cca_layer)).

## 1 Introduction

Cross-modality retrieval is the task of retrieving relevant items of a different modality than the search query (e.g., retrieving an image given a text query). One approach to tackle this problem is to define transformations which embed samples from different modalities into a common vector space. We can then project a query into this embedding space, and retrieve, using nearest neighbor search, a corresponding candidate projected from another modality.

A particularly successful class of models uses parametric nonlinear transformations (e.g., neural networks) for the embedding projections, optimized via a retrieval-specific objective such as a pairwise rank-

ing loss [15, 27]. This loss aims at decreasing the distance (a differentiable function such as Euclidean or cosine distance) between matching items, while increasing it between mismatching ones. Specialized extensions of this loss achieved state-of-the-art results in various domains such as natural language processing [10], image captioning [12], and text-to-image retrieval [29].

In a different approach, Yan and Mikolajczyk [31] propose to learn a joint embedding of text and images using Deep Canonical Correlation Analysis (DCCA) [2]. Instead of a pairwise ranking loss, DCCA directly optimizes the correlation of learned latent representations of the two views. Given the correlated embedding representations of the two views, it is possible to perform retrieval via cosine distance. The promising performance of their approach is also in line with the findings of Costa et al. [23] who state the following two hypotheses regarding the properties of efficient cross-modal retrieval spaces: First, the embedding spaces should account for low-level cross-modal correlations and second, they should enable semantic abstraction. In [31], both properties are met by a deep neural network — learning abstract representations — that is optimized with DCCA ensuring highly correlated latent representations.

In summary, the optimization of pairwise ranking losses yields embedding spaces that are useful for retrieval, and allows incorporating domain knowledge into the loss function. On the other hand, DCCA is designed to maximize correlation—which has already proven to be useful for cross-modality retrieval [31]—but does not allow to use loss formulations specialized for the task at hand.

In this paper, we propose a method to combine both approaches in a way that retains their advantages. We develop a *Canonical Correlation Analysis Layer* (CCAL) that can be inserted into a dual-view neural network to produce a maximally correlated embedding space for its latent representations. We can then apply task specific loss functions, in particular the pairwise ranking loss, on the output of

this layer. To train a network using the CCA layer we describe how to backpropagate the gradient of this loss function to the dual-view neural network while relying on automatic differentiation tools such as *Theano* [28] or *Tensorflow* [1]. In our experiments, we show that our proposed method performs better than DCCA and models using pairwise ranking loss alone, especially when little training data is available.

Figure 1 compares our proposed approach to the alternatives discussed above. DCCA defines an objective optimizing a dual-view neural network such that its two views will be maximally correlated (Figure 1a). Pairwise ranking losses are loss functions to optimize a dual-view neural network such that its two views are well-suited for nearest-neighbor retrieval in the embedding space (Figure 1b). In our approach, we boost optimization of a pairwise ranking loss based on cosine distance by placing a special-purpose layer, the CCA projection layer, between a dual-view neural network and the optimization target (Figure 1c). Our experiments in Section 5 will show the effectiveness of this proposal.

## 2 Canonical Correlation Analysis

In this section we review the concepts of CCA, the basis for our methodology. Let  $\mathbf{x} \in \mathbb{R}^{d_x}$  and  $\mathbf{y} \in \mathbb{R}^{d_y}$  denote two random column vectors with covariances  $\Sigma_{xx}$  and  $\Sigma_{yy}$  and cross-covariance  $\Sigma_{xy}$ . The objective of CCA is to find two matrices  $\mathbf{A}^* \in \mathbb{R}^{d_x \times k}$  and  $\mathbf{B}^* \in \mathbb{R}^{d_y \times k}$  composed of  $k$  paired column vectors  $\mathbf{A}_j$  and  $\mathbf{B}_j$  (with  $k \leq d_x$  and  $k \leq d_y$ ) that project  $\mathbf{x}$  and  $\mathbf{y}$  into a common space maximizing their componentwise correlation:

$$(\mathbf{A}^*, \mathbf{B}^*) = \arg \max_{\mathbf{A}, \mathbf{B}} \sum_{j=1}^k \text{corr}(\mathbf{A}'_j \mathbf{x}, \mathbf{B}'_j \mathbf{y}) \quad (1)$$

$$= \arg \max_{\mathbf{A}, \mathbf{B}} \sum_{j=1}^k \frac{\mathbf{A}'_j \Sigma_{xy} \mathbf{B}_j}{\sqrt{\mathbf{A}'_j \Sigma_{xx} \mathbf{A}_j \mathbf{B}'_j \Sigma_{yy} \mathbf{B}_j}} \quad (2)$$

Since the objective of CCA is invariant to scaling of the projection matrices, we constrain the projected dimensions to have unit variance. Furthermore, CCA seeks subsequently uncorrelated projection vectors, arriving at the equivalent formulation:

$$(\mathbf{A}^*, \mathbf{B}^*) = \arg \max_{\mathbf{A}' \Sigma_{xx} \mathbf{A} = \mathbf{B}' \Sigma_{yy} \mathbf{B} = \mathbf{I}_k} \text{tr}(\mathbf{A}' \Sigma_{xy} \mathbf{B}) \quad (3)$$

Let  $\mathbf{T} = \Sigma_{xx}^{-1/2} \Sigma_{xy} \Sigma_{yy}^{-1/2}$ , and let  $\mathbf{U} \text{diag}(\mathbf{d}) \mathbf{V}'$  be the Singular Value Decomposition (SVD) of  $\mathbf{T}$  with ordered singular values  $d_i \geq d_{i+1}$ . As shown in [19],

we obtain  $\mathbf{A}^*$  and  $\mathbf{B}^*$  from the top  $k$  left- and right-singular vectors of  $\mathbf{T}$ :

$$\mathbf{A}^* = \Sigma_{xx}^{-1/2} \mathbf{U}_{:k} \quad \mathbf{B}^* = \Sigma_{yy}^{-1/2} \mathbf{V}_{:k} \quad (4)$$

Moreover, the correlation in the projection space is the sum of the top  $k$  singular values:<sup>1</sup>

$$\text{corr}(\mathbf{A}^{*'} \mathbf{x}, \mathbf{B}^{*'} \mathbf{y}) = \sum_{i \leq k} d_i \quad (5)$$

In practice, the covariances and cross-covariance of  $\mathbf{x}$  and  $\mathbf{y}$  are usually not known, but estimated from a training set of  $m$  paired vectors, expressed as matrices  $\mathbf{X} \in \mathbb{R}^{d_x \times m}$ ,  $\mathbf{Y} \in \mathbb{R}^{d_y \times m}$  by:

$$\hat{\Sigma}_{xx} = \frac{1}{m-1} \bar{\mathbf{X}} \bar{\mathbf{X}}' + r \mathbf{I} \quad \text{and} \quad \hat{\Sigma}_{xy} = \frac{1}{m-1} \bar{\mathbf{X}} \bar{\mathbf{Y}}'. \quad (6)$$

$\bar{\mathbf{X}}$  is the centered version of  $\mathbf{X}$ .  $\hat{\Sigma}_{yy}$  is defined analogously to  $\hat{\Sigma}_{xx}$ . Additionally, we apply a regularization parameter  $r \mathbf{I}$  to ensure that the covariance matrices are positive definite. Substituting these estimates for  $\Sigma_{xx}$ ,  $\Sigma_{xy}$  and  $\Sigma_{yy}$ , respectively, we can compute  $\mathbf{A}^*$  and  $\mathbf{B}^*$  using Equation (4).

## 3 Cross-Modality Retrieval Baselines

In this section we review the two most related works forming the basis for our approach.

### 3.1 Deep Canonical Correlation Analysis

Andrew et al. [2] propose an extension of CCA to learn parametric nonlinear transformations of two random vectors, such that their correlation is maximized. Let  $\mathbf{a} \in \mathbb{R}^{d_a}$  and  $\mathbf{b} \in \mathbb{R}^{d_b}$  denote two random vectors, and let  $\mathbf{x} = f(\mathbf{a}; \Theta_f)$  and  $\mathbf{y} = g(\mathbf{b}; \Theta_g)$  denote their nonlinear transformations, parameterized by  $\Theta_f$  and  $\Theta_g$ . DCCA optimizes the parameters  $\Theta_f$  and  $\Theta_g$  to maximize the correlation of the topmost hidden representations  $\mathbf{x}$  and  $\mathbf{y}$ . For  $d_x = d_y = k$ , this objective corresponds to Equation 5, i.e., the sum of all singular values of  $\mathbf{T}$ , also called the trace norm:

$$\text{corr}(\mathbf{A}^{*'} f(\mathbf{a}; \Theta_f), \mathbf{B}^{*'} g(\mathbf{b}; \Theta_g)) = \|\mathbf{T}\|_{\text{tr}}. \quad (7)$$

Andrew et al. [2] show how to compute the gradient of this *Trace Norm Objective* (TNO) with respect to  $\mathbf{x}$  and  $\mathbf{y}$ . Assuming  $f$  and  $g$  are differentiable with respect to  $\Theta_f$  and  $\Theta_g$  (as is the case for neural networks), this allows to optimize the nonlinear transformations via gradient-based methods.

Yan and Mikolajczyk [31] suggest the following procedure to utilize DCCA for cross-modality retrieval:

<sup>1</sup>We understand the correlation of two vectors to be defined as  $\text{corr}(\mathbf{x}, \mathbf{y}) = \sum_i \sum_j \text{corr}(x_i, y_j)$ .

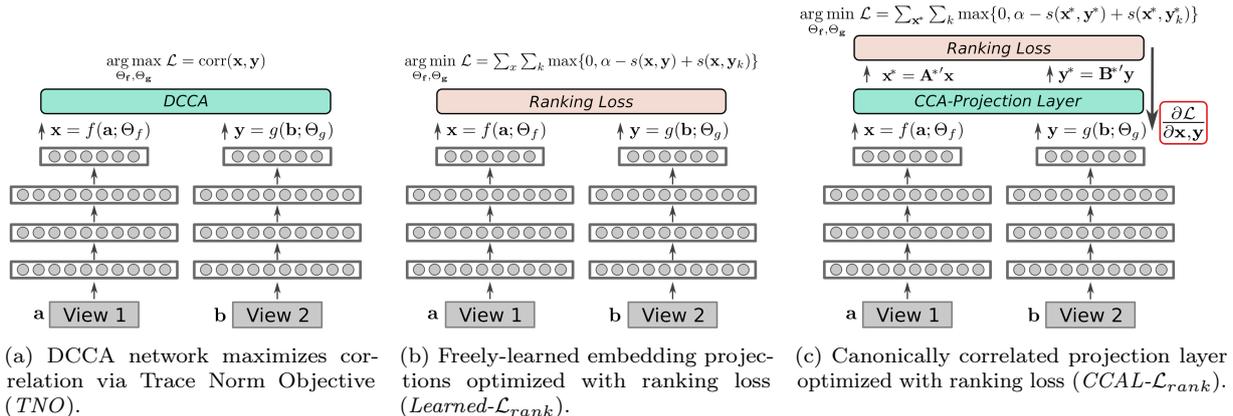


Figure 1: Sketches of cross-modality retrieval networks. The proposed model in (c) unifies (a) and (b) and takes advantage of both: componentwise correlated CCA projections and a pairwise ranking loss for cross-modality embedding space learning. We emphasize that our proposal in (c) requires to backpropagate the ranking loss  $\mathcal{L}$  through the analytical computation of the optimally correlated CCA embedding projections  $\mathbf{A}^*$  and  $\mathbf{B}^*$  (see Equation (4)). We thus need to compute their partial derivatives with respect to the network’s hidden representations  $\mathbf{x}$  and  $\mathbf{y}$ , i.e.  $\frac{\partial \mathbf{A}^*}{\partial \mathbf{x}, \mathbf{y}}$  and  $\frac{\partial \mathbf{B}^*}{\partial \mathbf{x}, \mathbf{y}}$  (addressed in Section 4).

first, neural networks  $f$  and  $g$  are trained using the TNO, with  $\mathbf{a}$  and  $\mathbf{b}$  representing different views of an entity (e.g. image and text); then, after the training is finished, the CCA projections are computed using Equation (4), and all retrieval candidates are projected into the embedding space; finally, at test time, queries of either modality are projected into the embedding space, and the best-matching sample from the other modality is found through nearest neighbor search using the cosine distance. Figure 2 provides a summary of the entire retrieval pipeline. In our experiments, we will refer to this approach as *DCCA-2015*.

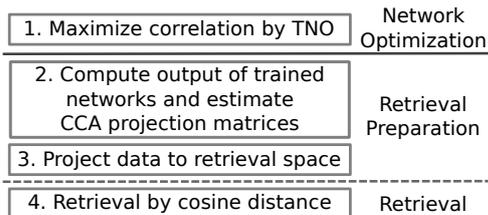


Figure 2: *DCCA* retrieval pipeline proposed in [31]. Note that all processing steps below the solid line are performed after network optimization is complete.

*DCCA* is limited by design to use the objective function described in Equation (7), and only seeks to maximize the correlation in the embedding space. During training, the CCA projection matrices are never computed, nor are the samples projected into the common retrieval space. All the retrieval steps—most importantly, the computation of CCA projections—are performed only once after the networks  $f$  and  $g$  have been optimized. This restricts potential applications, because we cannot use the projected data as an input

to subsequent layers or task-specific objectives. We will show how our approach overcomes this limitation in Section 4.

### 3.2 Pairwise Ranking Loss

Kiros et al. [15] learn a multi-modal joint embedding space for images and text. They use the cosine of the angle between two corresponding vectors  $\mathbf{x}$  and  $\mathbf{y}$  as a scoring function, i.e.,  $s(\mathbf{x}, \mathbf{y}) = \cos(\mathbf{x}, \mathbf{y})$ . Then, they optimize a pairwise ranking loss

$$\mathcal{L}_{rank} = \sum_{\mathbf{x}} \sum_k \max\{0, \alpha - s(\mathbf{x}, \mathbf{y}) + s(\mathbf{x}, \mathbf{y}_k)\} \quad (8)$$

where  $\mathbf{x}$  is an embedded sample of the first modality,  $\mathbf{y}$  is the matching embedded sample of the second modality, and  $\mathbf{y}_k$  are the contrastive (mismatching) embedded samples of the second modality (in practice, all mismatching samples in the current mini-batch). The hyper-parameter  $\alpha$  defines the margin of the loss function. This loss encourages an embedding space where the cosine distance between matching samples is lower than the cosine distance of mismatching samples.

In this setting, the networks  $f$  and  $g$  have to learn the embedding projections freely from randomly initialized weights. Since the projections are learned from scratch by optimizing a ranking loss, in our experiments we denote this approach by *Learned- $\mathcal{L}_{rank}$* . Figure 1b shows a sketch of this paradigm.

## 4 Learning with Canonically Correlated Embedding Projections

In the following we explain how to bring both concepts — DCCA and Pairwise Ranking Losses — together to enhance cross-modality embedding space learning.

### 4.1 Motivation

We start by providing an intuition on why we expect this combination to be fruitful: *DCCA-2015* maximizes the correlation between the latent representations of two different neural networks via the TNO derived from classic CCA. As correlation and cosine distance are related, we can also use such a network for cross-modality retrieval [31]. Kiros et al. [15], on the other hand, learn a cross-modality retrieval embedding by optimizing an objective customized for the task at hand. The motivation for our approach is that we want to benefit from both: a task specific retrieval objective, and componentwise optimally correlated embedding projections.

To achieve this, we devise a *CCA layer* that analytically computes the CCA projections  $\mathbf{A}^*$  and  $\mathbf{B}^*$  during training, and projects incoming samples into the embedding space. The projected samples can then be used in subsequent layers, or for computing task-specific losses such as the pairwise ranking loss. Figure 1c illustrates the central idea of our combined approach. Compared to Figure 1b, we insert an additional linear transformation. However, this transformation is not learned (otherwise it could be merged with the previous layer, which is not followed by a nonlinearity). Instead, it is computed to be the transformation that maximizes componentwise correlation between the two views.  $\mathbf{A}^*$  and  $\mathbf{B}^*$  in Figure 1c are the very projections given by Equation (4) in Section 2.

In theory, optimizing a pairwise ranking loss alone could yield projections equivalent to the ones computed by CCA. In practice, however, we observe that the proposed combination gives much better cross-modality retrieval results (see Section 5).

Our design requires backpropagating errors through the analytical computation of the CCA projection matrices. DCCA [2] does not cover this, since projecting the data is not necessary for optimizing the TNO. In the remainder of this section, we discuss how to establish gradient flow (backpropagation) through CCA’s optimal projection matrices. In particular, we require the partial derivatives  $\frac{\partial \mathbf{A}^*}{\partial \mathbf{x}, \mathbf{y}}$  and  $\frac{\partial \mathbf{B}^*}{\partial \mathbf{x}, \mathbf{y}}$  of the projections with respect to their input representations  $\mathbf{x}$  and  $\mathbf{y}$ . This will allow us to use CCA as a layer within a multi-modality neural network, instead of as a final objective (TNO) for correlation maximization only.

### 4.2 Gradient of CCA Projections

As mentioned above, we can compute the canonical correlation along with the optimal projection matrices from the singular value decomposition  $\mathbf{T} = \Sigma_{xx}^{-1/2} \Sigma_{xy} \Sigma_{yy}^{-1/2} = \mathbf{U} \text{diag}(\mathbf{d}) \mathbf{V}'$ . Specifically, we obtain the correlation as  $\text{corr}(\mathbf{A}^* \mathbf{x}, \mathbf{B}^* \mathbf{y}) = \sum_i d_i$ , and the projections as  $\mathbf{A}^* = \Sigma_{xx}^{-1/2} \mathbf{U}$  and  $\mathbf{B}^* = \Sigma_{yy}^{-1/2} \mathbf{V}$ . For DCCA, it suffices to compute the gradient of the total correlation wrt.  $\mathbf{x}$  and  $\mathbf{y}$  in order to backpropagate it through the two networks  $f$  and  $g$ . Using the chain rule, Andrew et al. decompose this into the gradients of the total correlation wrt.  $\Sigma_{xx}$ ,  $\Sigma_{xy}$  and  $\Sigma_{yy}$ , and the gradients of those wrt.  $\mathbf{x}$  and  $\mathbf{y}$  [2]. Their derivations of the former make use of the fact that both the gradient of  $\sum_i d_i$  wrt.  $\mathbf{T}$  and the gradient of  $\|\mathbf{T}\|_{\text{tr}}$  (the trace norm objective in Equation (7)) wrt.  $\mathbf{T}'\mathbf{T}$  have a simple form; see Section 7 in [2] for details.

In our case where we would like to backpropagate errors through the CCA transformations, we instead need the gradients of the projected data  $\mathbf{x}^* = \mathbf{A}^* \mathbf{x}$  and  $\mathbf{y}^* = \mathbf{B}^* \mathbf{y}$  wrt.  $\mathbf{x}$  and  $\mathbf{y}$ , which requires the partial derivatives  $\frac{\partial \mathbf{A}^*}{\partial \mathbf{x}, \mathbf{y}}$  and  $\frac{\partial \mathbf{B}^*}{\partial \mathbf{x}, \mathbf{y}}$ . We could again decompose this into the gradients wrt.  $\mathbf{T}$ , the gradients of  $\mathbf{T}$  wrt.  $\Sigma_{xx}$ ,  $\Sigma_{xy}$  and  $\Sigma_{yy}$  and the gradients of those wrt.  $\mathbf{x}$  and  $\mathbf{y}$ . However, while the gradients of  $\mathbf{U}$  and  $\mathbf{V}$  wrt.  $\mathbf{T}$  are known [22], they involve solving  $O((d_x d_y)^2)$  linear  $2 \times 2$  systems. Instead, we reformulate the solution to use two symmetric eigendecompositions  $\mathbf{T}\mathbf{T}' = \mathbf{U} \text{diag}(\mathbf{e}) \mathbf{U}'$  and  $\mathbf{T}'\mathbf{T} = \mathbf{V} \text{diag}(\mathbf{e}) \mathbf{V}'$  (Equation 270 in [24]). This gives us the same left and right eigenvectors we would obtain from the SVD, along with the squared singular values ( $e_i = d_i^2$ ). The gradients of eigenvectors of symmetric real eigensystems have a simple form [17] and both  $\mathbf{T}\mathbf{T}'$  and  $\mathbf{T}'\mathbf{T}$  are differentiable wrt.  $\mathbf{x}$  and  $\mathbf{y}$ .

To summarize: in order to obtain an efficiently computable definition of the gradient for CCA projections, we have reformulated the forward pass (the computation of the CCA transformations). Our formulation using two eigen-decompositions translates into a series of computation steps that are differentiable in a graph-based, auto-differentiating math compiler such as *Theano* [28], which, together with the chain rule, gives an efficient implementation of the CCA layer gradient for training our network<sup>2</sup>. For a detailed description of the CCA layer forward pass we refer to Algorithm 1 in the Appendix of this article. As the technical implementation is not straight-forward, we also discuss the crucial steps in the Appendix.

Thus, we now have the means to benefit from the optimal CCA projections but still optimize for a task-specific objective. In particular, we utilize the *pairwise ranking loss* of Equation (8) on top of an in-

<sup>2</sup>The code of our implementation of the CCA layer is available at [https://github.com/CPJKU/cca\\_layer](https://github.com/CPJKU/cca_layer).

intermediate CCA embedding projection layer. We denote the proposed retrieval network of Figure 1c as  $CCAL-\mathcal{L}_{rank}$  in our experiments ( $CCAL$  refers to CCA Layer).

## 5 Experiments

We evaluate our approach ( $CCAL-\mathcal{L}_{rank}$ ) in cross-modality retrieval experiments on two image-to-text and one audio-to-sheet-music dataset. Additionally, we provide results on two zero-shot text-to-image retrieval scenarios proposed in [25]. For comparison, we consider the approach of [31] ( $DCCA-2015$ ), our own implementation of the TNO (denoted by  $DCCA$ ), as well as the freely learned projection embeddings ( $Learned-\mathcal{L}_{rank}$ ) optimizing the ranking loss of [15].

The task for all three datasets is to retrieve the correct counterpart when given an instance of the other modality as a search query. For retrieval, we use the cosine distance in embedding space for all approaches. First, we embed all candidate samples of the target modality into the retrieval embedding space. Then, we embed the query element  $\mathbf{y}$  with the second network and select its nearest neighbor  $\mathbf{x}_j$  of the target modality. Fig. 3 shows a sketch of this retrieval by embedding space learning paradigm.

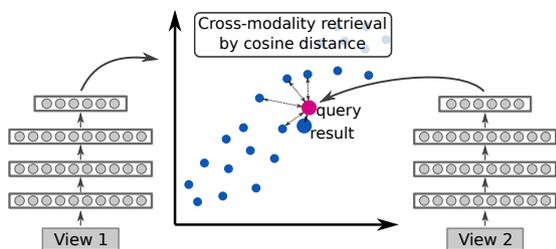


Figure 3: Sketch of cross-modality retrieval. The blue dots are the embedded candidate samples. The red dot is the embedding of the search query. The larger blue dot highlights the closest candidate selected as the retrieval result.

As evaluation measures, we consider the  $Recall@k$  ( $R@k$  in %) as well as the  $Median Rank$  ( $MR$ ) and the  $Mean Reciprocal Rank$  ( $MRR$  in %). The  $R@k$  rate (higher is better) is the ratio of queries which have the correct corresponding counterpart in the first  $k$  retrieval results. The  $MR$  (lower is better) is the median position of the target in a similarity-ordered list of available candidates. Finally, we define the  $MRR$  (higher is better) as the mean value of  $1/rank$  over all queries where  $rank$  is again the position of the target in the similarity ordered list of available candidates.

### 5.1 Image-Text Retrieval

In the first part of our experiments, we consider  $Flickr30k$  and  $IAPR TC-12$ , two publicly available datasets for image-text cross-modality retrieval. Flickr30k consists of image-caption pairs, where each image is annotated with five different textual descriptions. The train-validation-test split for Flickr30k is 28000-1000-1000. In terms of evaluation setup, we follow *Protocol 3* of [31] and concatenate the five available captions into one, meaning that only one, but richer text annotation remains per image. This is done for all three sets of the split. The second image-text dataset, IAPR TC-12, contains 20000 natural images where only one—but compared to Flickr30k more detailed—caption is available for each image. As no predefined train-validation-test split is provided, we randomly select 1000 images for validation and 2000 for testing, and keep the rest for training. [31] also use 2000 images for testing, but did not explicitly mention holdout images for validation. Table 1 shows an example image along with its corresponding captions or caption for either dataset.

The input to our networks is a 4096-dimensional image feature vector along with a corresponding text vector representation which has dimensionality 5793 for Flickr30k and 2048 for IAPR TC-12. The image embedding is computed from the last hidden layer of a network pretrained on ImageNet [7] (layer  $fc7$  of  $CNN\_S$  by [4]). In terms of text pre-processing, we follow [31], tokenizing and lemmatizing the raw captions as the first step. Based on the lemmatized captions, we compute  $l2$ -normalized TF/IDF-vectors, omitting words with an overall occurrence smaller than five for Flickr30k and three for IAPR TC-12, respectively. The image representation is processed by a linear dense layer with 128 units, which will also be the dimensionality  $k$  of the resulting retrieval embedding. The text vector is fed through two batch-normalized [11] dense layers of 1024 units each and the ELU activation function [6]. As a last layer for the text representation network, we again apply a dense layer with 128 linear units.

For a fair comparison, we keep the structure and number of parameters of all networks in our experiments the same. The only difference between the networks are the objectives and the hyper-parameters used for optimization. Optimization is performed using Stochastic Gradient Descent (SGD) with the *adam* update rule [14] (for details please see our appendix).

Table 2 lists our results on IAPR TC-12. Along with our experiments, we also show the results reported in [31] as a reference ( $DCCA-2015$ ). However, a direct comparison to our results may not be fair:  $DCCA-2015$  uses a different ImageNet-pretrained network for the image representation, and finetunes this network while we keep it fixed. This is because our interest is in comparing the methods in a stable set-



A man in a white cowboy hat reclines in front of a window in an airport.

A young man rests on an airport seat with a cowboy hat over his face.

A woman relaxes on a couch, with a white cowboy hat over her head.

A man is sleeping inside on a bench with his hat over his eyes.

A person is sleeping at an airport with a hat on their head.



a green and brown embankment with brown houses on the right and a light brown sandy beach at the dark blue sea on the left; a dark mountain range behind it and white clouds in a light blue sky in the background;

Table 1: Example images for Flickr30k (top) and IAPR TC-12 (bottom)

ting, not in obtaining the best possible results. Our implementation of the TNO (*DCCA*) uses the same objective as *DCCA-2015*, but is trained using the same network architecture as our remaining models and permits a direct comparison. Additionally, we repeat each of the experiments 10 times with different initializations and report the mean for each of the evaluation measures.

When taking a closer look at Table 2, we observe that our results achieved by optimizing the TNO (*DCCA*) surpass the results reported in [31]. We already discussed above that the two versions are not directly comparable. However, given this result, we consider our implementation of *DCCA* as a valid baseline for our experiments in Section 5.2 where no results are available in the literature. When looking at the performance of *CCAL-L<sub>rank</sub>* we further observe that it outperforms all other methods, although the difference to *DCCA* is not pronounced for all of the measures. Comparing *CCAL-L<sub>rank</sub>* with the freely-learned projection matrices (*Learned-L<sub>rank</sub>*) we observe a much larger performance gap. This is interesting, as in principle the learned projections could converge to exactly the same solution as *CCAL-L<sub>rank</sub>*. We take this as a quantitative confirmation that the learning process benefits from CCA’s optimal projection matrices.

In Table 3, we list our results on the Flickr30k dataset. As above, we show the retrieval performances of [31] as a baseline along with our results and observe similar behavior as on IAPR TC-12. Again, we point out the poor performance of the freely-learned projections (*Learned-L<sub>rank</sub>*) in this experiment. Keeping this observation in mind, we will notice a different behavior in the experiments in Section 5.2.

Note that there are various other methods reporting results on Flickr30k [13, 27, 18, 15] which partly surpass ours, for example by using more elaborate processing of the textual descriptions or more powerful ImageNet models. We omit these results as we focus on the comparison of *DCCA* and freely-learned projections with the proposed CCA embedding layer.

## 5.2 Audio-Sheet-Music Retrieval

For the second set of experiments, we consider the Nottingham piano midi dataset [3]. The dataset is a collection of midi files split into train, validation and test set already used by [8] for experiments on end-to-end score-following in sheet-music images. Here, we tackle the problem of audio-sheet-music retrieval, i.e. matching short snippets of music (audio) to corresponding parts in the sheet music (image). Figure 4 shows examples of such correspondences.

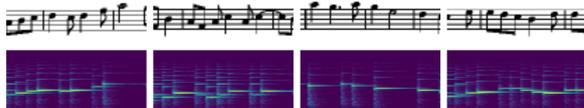


Figure 4: Example of the data considered for audio-sheet-music (image) retrieval. Top: short snippets of sheet music images. Bottom: Spectrogram excerpts of the corresponding music audio.

We conduct this experiment for two reasons: First, to show the advantage of the proposed method over different domains. Second, the data and application is of high practical relevance in the domain of Music Information Retrieval (MIR). A system capable of linking sheet music (images) and the corresponding music (audio) would be useful in many content-based musical retrieval scenarios.

In terms of audio preparation, we compute log frequency spectrograms with a sample rate of 22.05kHz, a FFT window size of 2048, and a computation rate of 31.25 frames per second. These spectrograms (136 frequency bins) are then directly fed into the audio part of the cross-modality networks. Figure 4 shows a set of audio-to-sheet correspondences presented to our network for training. One audio excerpt comprises 100 frames and the dimension of the sheet image snippet is  $40 \times 100$  pixels. Overall this results in 270,705

Table 2: Retrieval results on IAPR TC-12. “DCCA-2015” is taken from [31].

Method	Image-to-Text					Text-to-Image				
	R@1	R@5	R@10	MR	MRR	R@1	R@5	R@10	MR	MRR
DCCA-2015	30.2	57.0	-	-	42.6	29.5	60.0	-	-	41.5
DCCA	31.0	58.7	70.4	3.6	43.9	29.5	58.2	70.5	4.0	42.7
Learned- $\mathcal{L}_{rank}$	22.3	50.7	63.8	5.2	35.7	21.6	50.1	63.3	5.5	35.1
CCAL- $\mathcal{L}_{rank}$	31.6	61.0	72.2	3.0	45.0	29.6	60.0	72.2	3.6	43.5

Table 3: Retrieval results on Flickr30k. “DCCA-2015” is taken from [31].

Method	Image-to-Text					Text-to-Image				
	R@1	R@5	R@10	MR	MRR	R@1	R@5	R@10	MR	MRR
DCCA-2015	27.9	56.9	68.2	4	-	26.8	52.9	66.9	4	-
DCCA	31.6	59.2	69.3	3.3	44.2	30.3	58.3	69.2	3.8	43.1
Learned- $\mathcal{L}_{rank}$	23.7	50.5	63.0	5.3	36.3	23.6	51.0	62.5	5.2	36.5
CCAL- $\mathcal{L}_{rank}$	32.0	59.2	70.4	3.2	44.8	29.9	58.8	70.2	3.7	43.3

train, 18,046 validation and 16,042 test audio-sheet-music pairs. This is an order of magnitude more training data than for the image-to-text datasets of the previous section.

In the experiments in Section 5.1, we relied on pre-trained ImageNet features and relatively shallow fully connected text-feature processing networks. The model here differs from this, as it consists of two deep convolutional networks learned entirely from scratch. Our architecture is a VGG-style [26] network consisting of sequences of  $3 \times 3$  convolution stacks followed by  $2 \times 2$  max pooling. To reduce the dimensionality to the desired correlation space dimensionality  $k$  (in this case 32) we insert as a final building block a  $1 \times 1$  convolution having  $k$  feature maps followed by global average pooling [16] (for further architectural details we again refer to the appendix of this manuscript).

Table 4 lists our result on audio-to-sheet music retrieval. As in the experiments on images and text, the proposed CCA projection embedding layer trained with pairwise ranking loss outperforms the other models. Recalling the results from Section 5.1, we observe an increased performance of the freely-learned embedding projections. On measures such as R@5 or R@10 it achieves similar to or better performance than *DCCA*. One of the reasons for this could be the fact that there is an order of magnitude more training data available for this task to learn the projection embedding from random initialization. Still, our proposed combination of both concepts (*CCAL- $\mathcal{L}_{rank}$* ) achieves highest retrieval scores.

### 5.3 Performance in Small Data Regime

The above results suggest that the benefit of using a CCA projection layer (*CCAL- $\mathcal{L}_{rank}$* ) over a freely-

learned projection becomes most evident when few training data is available. To examine this assumption, we repeat the audio-to-sheet-music experiment of the previous section, but use only 10% of the original training data ( $\approx 27000$  samples). We stress the fact that the learned embedding projection of *Learned- $\mathcal{L}_{rank}$*  could converge to exactly the same solution as the CCA projections of *CCAL- $\mathcal{L}_{rank}$* . Table 5 summarizes the low data regime results for the three methods. Consistent with our hypothesis, we observe a larger gap between *Learned- $\mathcal{L}_{rank}$*  and *CCAL- $\mathcal{L}_{rank}$*  compared to the one obtained with all training data in Table 4. We conclude that a network might be able to learn suitable embedding projections when sufficient training data is available. However, when having fewer training samples, the proposed CCA projection layer strongly supports embedding space learning. In addition, we also looked into the retrieval performance of *Learned- $\mathcal{L}_{rank}$*  and *CCAL- $\mathcal{L}_{rank}$*  on the training set and observe comparable performance. This indicates that the CCA layer also acts as a regularizer and helps to generalize to unseen samples.

### 5.4 Zero-Shot Image-Text Retrieval

Our last set of experiments focuses on a slightly modified retrieval setting, namely image-text *zero-shot retrieval* [25]. Given a set of image-text pairs originating from  $C$  different categories the data is split into a class-disjoint training, validation and test sets having no categorical overlap. This implies that at test time we aim to retrieve images from textual queries describing categories (semantic concepts) never seen before, neither for training, nor for validation.

Reed et al. [25] collected and provided textual descriptions for two publicly available datasets, the

Table 4: Retrieval results on Nottingham dataset (Audio-to-Sheet-Music Retrieval).

Method	Sheet-to-Audio					Audio-to-Sheet				
	R@1	R@5	R@10	MR	MRR	R@1	R@5	R@10	MR	MRR
DCCA	42.0	88.2	93.3	2	62.2	44.6	87.9	93.2	2	63.5
Learned- $\mathcal{L}_{rank}$	40.7	89.6	95.6	2	61.7	41.4	88.9	95.4	2	61.9
CCAL- $\mathcal{L}_{rank}$	44.1	93.3	97.7	2	65.3	44.5	91.6	96.7	2	64.9

Table 5: Retrieval results on audio-to-sheet-music retrieval when using only 10% of the train data.

Method	Sheet-to-Audio					Audio-to-Sheet				
	R@1	R@5	R@10	MR	MRR	R@1	R@5	R@10	MR	MRR
DCCA	20.0	53.6	65.4	5	35.3	22.7	54.7	65.8	4	37.3
Learned- $\mathcal{L}_{rank}$	11.3	35.2	47.6	12	23.0	12.6	35.2	47.2	12	23.7
CCAL- $\mathcal{L}_{rank}$	22.2	59.2	70.7	4	38.8	25.0	59.3	70.9	4	40.4

CUB-200 bird image dataset [30] and the Oxford Flowers dataset [21]. According to the definition of zero-shot retrieval above we follow [25] and split CUB into 100 train, 50 validation and 50 test categories. Flowers is split into 82 train and 20 validation / test classes respectively. Figure 5 shows some example images along with their textual descriptions.

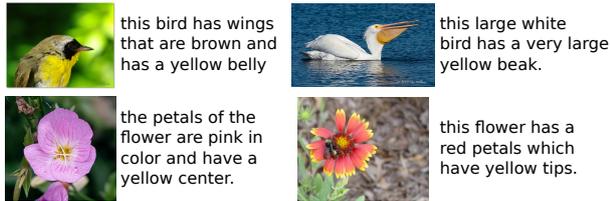


Figure 5: Example images of CUB-200 birds and Oxford Flowers along with textual descriptions collected by Reed et al. [25] for zero-shot retrieval from text.

Besides the modified, harder retrieval setting there is a second difference to the text-image retrieval experiments carried out in Section 5.1. Instead of using hand engineered textual features (e.g. TF-IDF) or unsupervised textual feature learning (e.g. word2vec [20]) the authors in [25] employ Convolutional Recurrent Neural Networks (CRNN) to learn the latent text representations directly from the raw descriptions. In particular, they feed the descriptions as one-hot-word encodings to the text processing part of their networks. In terms of image representations they still rely on 1024-dimensional pretrained ImageNet features. The feature learning part and the network architectures used for our experiments follows exactly the descriptions provided in [25]. The sole difference is, that we again replace the topmost embedding layer with the proposed CCA projection layer in combination with a pairwise ranking loss.

Table 6 compares the retrieval results of the respective methods on the two zero-shot retrieval datasets.

To allow for a direct comparison with the results reported in [25] we follow their evaluation setup and report the *Average Precision (AP@50)*. The AP@50 is the percentage of the top-50 scoring images whose class matches that of the text query, averaged over the 50 test classes. In [25] the best retrieval performance for both datasets (when considering only feature learning) is achieved by having a CRNN directly processing the textual descriptions. What is also interesting is the substantial performance gain with respect to unsupervised word2vec features.

For the Birds dataset, as an alternative to the textual descriptions, there are manually created fine-grained attributes available for each of the images. When relying on these attributes Reed et al. report state of the art results on the dataset [25] not reached by their text processing neural networks.

In the bottom part of Table 6, we report the performance of the same architectures optimized using our proposed CCA layer in combination with a pairwise ranking loss. We observe that the CCA layer is able to improve the performance of both models on both datasets. The gain in retrieval performance within a model class is largest for the convolution only (CNN) text processing models ( $\approx 9$  percentage points for the Flowers dataset and  $\approx 6$  for CUB). For the birds dataset the *Word CNN + CCAL* even outperforms the models relying on manually encoded attributes by achieving an AP@50 of 52.2.

## 6 Discussion and Conclusion

We have shown how to use the optimal projection matrices of CCA as the weights of an embedding layer within a multi-view neural network. With this CCA layer, it becomes possible to optimize for a specialized loss function (e.g., related to a retrieval task) on top of this, exploiting the correlation properties of a latent

Table 6: Zero-shot retrieval results on Cub and Flowers.

Method	Flowers	Birds
Attributes [25]	-	50.0
Word2Vec [25]	52.1	33.5
Word CNN [25]	56.3	43.3
Word CNN-RNN [25]	59.6	48.7
Word CNN + CCAL	62.2	52.2
Word CNN-RNN + CCAL	64.0	49.8

space provided by CCA. As this requires to establish gradient flow through CCA, we formulate it to allow easy computation of the partial derivatives  $\frac{\partial \mathbf{A}^*}{\partial \mathbf{x}, \mathbf{y}}$  and  $\frac{\partial \mathbf{B}^*}{\partial \mathbf{x}, \mathbf{y}}$  of CCA’s projection matrices  $\mathbf{A}^*$  and  $\mathbf{B}^*$  with respect to the input data  $\mathbf{x}$  and  $\mathbf{y}$ . With this formulation, we can incorporate CCA as a building block within multi-modality neural networks that produces maximally-correlated projections of its inputs. In our experiments, we use this building block within a cross-modality retrieval setting, optimizing a network to minimize a cosine distance based pairwise ranking loss of the componentwise-correlated CCA projections. Experimental results show that when using the cosine distance for retrieval (as is common for correlated views), this is superior to optimizing a network for maximally-correlated projections (as done in DCCA), or not using CCA at all. This observation holds in our experiments on a variety of different modality pairs as well as two different retrieval scenarios.

When investigating the experimental results in more detail, we find that the correlation-based methods (DCCA, CCAL) consistently outperform the models that learn the embedding projections from scratch. A direct comparison of DCCA with the proposed CCAL- $\mathcal{L}_{rank}$  reveals two learning scenarios where CCAL- $\mathcal{L}_{rank}$  is superior: (1) the low data regime, where we found that the CCA layer acts as a strong regularizer to prevent over-fitting; (2) when learning the entire retrieval representation (network parameterization) from scratch, not relying on pre-trained or hand-crafted features (see Section 5.2). Our intuition on this is that incorporating the task-specific retrieval objective already during training encourages the networks to learn embedding representations that are beneficial for retrieval at test-time. This is the important conceptual difference compared to the Trace Norm Objective (TNO) of DCCA, which does not focus on the retrieval task. However, when using the CCA layer we also inherit one drawback of the pairwise ranking loss, which is the additional hyperparameter (margin  $\alpha$ ) that needs to be determined on the validation set.

Finally, we would like to emphasize that our CCA layer is a general network component which could

provide a useful basis for further research, e.g., as an intermediate processing step for learning binary cross-modality retrieval representations.

## Acknowledgements

This work is supported by the Austrian Ministries BMVIT and BMWFW, and the Province of Upper Austria via the COMET Center SCCH, and by the European Research Council (ERC) under the EU’s Horizon 2020 Framework Programme (ERC Grant Agreement number 670035, project "Con Espressione"). The Tesla K40 used for this research was donated by the NVIDIA corporation.

## References

- [1] Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, et al. Tensorflow: Large-scale machine learning on heterogeneous distributed systems. *arXiv preprint arXiv:1603.04467*, 2016.
- [2] Galen Andrew, Raman Arora, Jeff Bilmes, and Karen Livescu. Deep canonical correlation analysis. In *Proceedings of the International Conference on Machine Learning*, pages 1247–1255, 2013.
- [3] Nicolas Boulanger-Lewandowski, Yoshua Bengio, and Pascal Vincent. Modeling temporal dependencies in high-dimensional sequences: Application to polyphonic music generation and transcription. In *Proceedings of the 29th International Conference on Machine Learning (ICML-12)*, pages 1159–1166, 2012.
- [4] K. Chatfield, K. Simonyan, A. Vedaldi, and A. Zisserman. Return of the devil in the details: Delving deep into convolutional nets. In *British Machine Vision Conference*, 2014.
- [5] Junyoung Chung, Çağlar Gülçehre, KyungHyun Cho, and Yoshua Bengio. Empirical evaluation of gated recurrent neural networks on sequence modeling. *CoRR*, abs/1412.3555, 2014.
- [6] Djork-Arné Clevert, Thomas Unterthiner, and Sepp Hochreiter. Fast and accurate deep network learning by exponential linear units (elus). *International Conference on Learning Representations (ICLR) (arXiv:1511.07289)*, 2015.
- [7] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. ImageNet: A Large-Scale Hierarchical Image Database. In *CVPR09*, 2009.

- [8] Matthias Dorfer, Andreas Arzt, and Gerhard Widmer. Towards score following in sheet music images. In *Proceedings of the International Society for Music Information Retrieval Conference (ISMIR)*, 2016.
- [9] David R Hardoon, Sandor Szedmak, and John Shawe-Taylor. Canonical correlation analysis: An overview with application to learning methods. *Neural computation*, 16(12):2639–2664, 2004.
- [10] Karl Moritz Hermann and Phil Blunsom. Multilingual distributed representations without word alignment. *arXiv preprint arXiv:1312.6173*, 2013.
- [11] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *CoRR*, abs/1502.03167, 2015.
- [12] Andrej Karpathy and Li Fei-Fei. Deep visual-semantic alignments for generating image descriptions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3128–3137, 2015.
- [13] Andrej Karpathy, Armand Joulin, and Fei Fei Li. Deep fragment embeddings for bidirectional image sentence mapping. In *Advances in neural information processing systems*, pages 1889–1897, 2014.
- [14] Diederik Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [15] Ryan Kiros, Ruslan Salakhutdinov, and Richard S Zemel. Unifying visual-semantic embeddings with multimodal neural language models. *arXiv preprint arXiv:1411.2539*, 2014.
- [16] Min Lin, Qiang Chen, and Shuicheng Yan. Network in network. *CoRR*, abs/1312.4400, 2013.
- [17] Jan R. Magnus. On differentiating eigenvalues and eigenvectors. *Econometric Theory*, 1(2):179–191, 1985.
- [18] Junhua Mao, Wei Xu, Yi Yang, Jiang Wang, and Alan L Yuille. Explain images with multimodal recurrent neural networks. *arXiv preprint arXiv:1410.1090*, 2014.
- [19] K.V. Mardia, J.T. Kent, and J.M. Bibby. *Multivariate analysis*. Probability and mathematical statistics. Academic Press, 1979.
- [20] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*, pages 3111–3119, 2013.
- [21] M-E. Nilsback and A. Zisserman. Automated flower classification over a large number of classes. In *Proceedings of the Indian Conference on Computer Vision, Graphics and Image Processing*, Dec 2008.
- [22] Théodore Papadopoulos and Manolis I.A. Lourakis. Estimating the Jacobian of the Singular Value Decomposition: Theory and Applications. In *Proceedings of the 6th European Conference on Computer Vision (ECCV)*, 2000.
- [23] Jose Costa Pereira, Emanuele Coviello, Gabriel Doyle, Nikhil Rasiwasia, Gert RG Lanckriet, Roger Levy, and Nuno Vasconcelos. On the role of correlation and abstraction in cross-modal multimedia retrieval. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 36(3):521–535, 2014.
- [24] K. B. Petersen and M. S. Pedersen. The matrix cookbook, nov 2012. Version 20121115.
- [25] Scott Reed, Zeynep Akata, Bernt Schiele, and Honglak Lee. Deep visual-semantic alignments for generating image descriptions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016.
- [26] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- [27] Richard Socher, Andrej Karpathy, Quoc V Le, Christopher D Manning, and Andrew Y Ng. Grounded compositional semantics for finding and describing images with sentences. *Transactions of the Association for Computational Linguistics*, 2:207–218, 2014.
- [28] Theano Development Team. Theano: A Python framework for fast computation of mathematical expressions. *arXiv e-prints*, abs/1605.02688, May 2016.
- [29] Ivan Vendrov, Ryan Kiros, Sanja Fidler, and Raquel Urtasun. Order-embeddings of images and language. *CoRR*, abs/1511.06361, 2016.
- [30] P. Welinder, S. Branson, T. Mita, C. Wah, F. Schroff, S. Belongie, and P. Perona. Caltech-UCSD Birds 200. Technical Report CNS-TR-2010-001, California Institute of Technology, 2010.
- [31] Fei Yan and Krystian Mikolajczyk. Deep correlation for matching images and text. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3441–3450, 2015.

# Appendix

## Implementation Details

Backpropagating the errors through the CCA projection matrices is not trivial. The optimal CCA projection matrices are given by  $\mathbf{A}^* = \Sigma_{xx}^{-1/2} \mathbf{U}$  and  $\mathbf{B}^* = \Sigma_{yy}^{-1/2} \mathbf{V}$ , where  $\mathbf{U}$  and  $\mathbf{V}$  are derived from the singular value decomposition of  $\mathbf{T} = \Sigma_{xx}^{-1/2} \Sigma_{xy} \Sigma_{yy}^{-1/2} = \mathbf{U} \text{diag}(\mathbf{d}) \mathbf{V}'$  (see Section 2). The proposed model needs to backpropagate the errors through the CCA transformations, i.e., it requires the gradients of the projected data  $\mathbf{x}^* = \mathbf{A}^* \mathbf{x}$  and  $\mathbf{y}^* = \mathbf{B}^* \mathbf{y}$  wrt.  $\mathbf{x}$  and  $\mathbf{y}$ . Applying the chain rule, this further requires the gradients of  $\mathbf{U}$  and  $\mathbf{V}$  wrt.  $\mathbf{T}$ , and the gradients of  $\mathbf{T}$ ,  $\Sigma_{xx}^{-1/2}$ ,  $\Sigma_{xy}$  and  $\Sigma_{yy}^{-1/2}$  wrt.  $\mathbf{x}$  and  $\mathbf{y}$ .

The main technical challenge is that common auto-differentiation tools such as Theano [28] or Tensor Flow [1] do not provide derivatives for the inverse squared root and singular value decomposition of a matrix.<sup>3</sup> To overcome this, we replace the inverse squared root of a matrix by using its Cholesky decomposition as described in [9]. Furthermore, we note that the singular value decomposition is required to obtain the matrices  $\mathbf{U}$  and  $\mathbf{V}$ , but in fact those matrices can alternatively be obtained by solving the eigendecomposition of  $\mathbf{T}\mathbf{T}' = \mathbf{U} \text{diag}(\mathbf{e}) \mathbf{U}'$  and  $\mathbf{T}'\mathbf{T} = \mathbf{V} \text{diag}(\mathbf{e}) \mathbf{V}'$  [24, Eq. 270]. This yields the same left and right eigenvectors we would obtain from the SVD (except for possibly flipped signs, which are easy to fix), along with the squared singular values ( $e_i = d_i^2$ ). Note that  $\mathbf{T}\mathbf{T}'$  and  $\mathbf{T}'\mathbf{T}$  are symmetric, and that the gradients of eigenvectors of symmetric real eigensystems have a simple form [17, Eq. 7]. Furthermore,  $\mathbf{T}\mathbf{T}'$  and  $\mathbf{T}'\mathbf{T}$  are differentiable wrt.  $\mathbf{x}$  and  $\mathbf{y}$ , enabling a sufficiently efficient implementation in a graph-based, auto-differentiating math compiler<sup>4</sup>.

The following section provides a detailed description of the implementation of the CCA layer.

## Forward Pass of CCA Projection Layer

For easier reproducibility, we provide a detailed description of the *forward pass* of the proposed CCA layer in Algorithm 1. To train the model, we need to propagate the gradient through the CCA layer (backward pass). We rely on auto-differentiation tools (in particular, *Theano*) implementing the gradient for each individual computation step in the forward pass, and connecting them using the chain rule.

The layer itself takes the latent feature representations (a batch of  $m$  paired vectors  $\mathbf{X} \in \mathbb{R}^{d_x \times m}$  and  $\mathbf{Y} \in \mathbb{R}^{d_y \times m}$ ) of the two network pathways  $f$  and  $g$

<sup>3</sup>Note that this is not relevant for the DCCA model introduced in [2] because it only derives the CCA projections after optimizing the TNO.

<sup>4</sup>The code of our implementation of the CCA layer is available at [https://github.com/CPJKU/cca\\_layer](https://github.com/CPJKU/cca_layer)

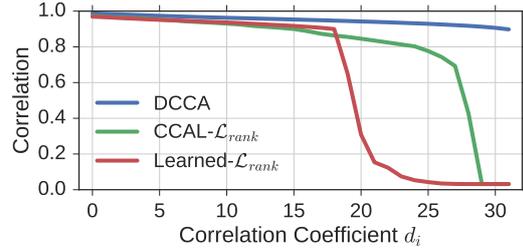


Figure 6: Comparison of the 32 correlation coefficients  $d_i$  (the dimensionality of the retrieval space is 32) of the topmost hidden representations  $\mathbf{x}$  and  $\mathbf{y}$  of the audio-to-sheet-music dataset and the respective optimization paradigm. The maximum correlation possible is 1.0 for each coefficient

as input and projects them with CCA’s analytical projection matrices. At train time, the layer uses the optimal projections computed from the current batch. When applying the layer at test time it uses the statistics and projections remembered from last training batch (which can of course be recomputed on a larger training batch to get more stable estimate).

As not all of the computation steps are obvious, we provide further details for the crucial ones. In line 12 and 13, we compute the Cholesky factorization instead of the matrix square root, as the latter has no gradients implemented in *Theano*. As a consequence, we need to transpose  $\mathbf{C}_{yy}^{-1}$  when computing  $\mathbf{T}$  in line 14 [9]. In line 15 and 16, we compute two eigen decompositions instead of one singular value decomposition (which also has no gradients implemented in *Theano*). In line 19, we flip the signs of first projection matrix to match the second to only have positive correlations. This property is required for retrieval with cosine distance. Finally, in line 24 and 25, the two views get projected using  $\mathbf{A}^*$  and  $\mathbf{B}^*$ . At test time we apply the projections computed and stored during training (line 17).

## Investigations on Correlation Structure

As an additional experiment we investigate the correlation structure of the learned representations for all three paradigms. For that purpose we compute the topmost hidden representation  $\mathbf{x}$  and  $\mathbf{y}$  of the audio-sheet-music-pairs and estimate the canonical correlation coefficients  $d_i$  of the respective embedding spaces. For the present example this yields 32 coefficients which is the dimensionality  $k$  of our retrieval embedding space. Figure 6 compares the correlation coefficients where 1.0 is the maximum value reachable. The most prominent observation in Figure 6 is the high correlation coefficients of the representation learned with *DCCA*. This structure is expected as the TNO focuses solely on correlation maximization. However, when recalling the results of Table 4 we

---

**Algorithm 1** Forward Pass of CCA Projection Layer.

---

```
1: Input of layer:  $\mathbf{X} \in \mathbb{R}^{d_x \times m}$  and  $\mathbf{Y} \in \mathbb{R}^{d_y \times m}$  ▷ hidden representation of current batch
2: Returns:  $\mathbf{X}^*$  and  $\mathbf{Y}^*$  ▷ CCA projected hidden representation
3: Parameters of layer:  $\mu_x, \mu_y$  and  $\mathbf{A}^*, \mathbf{B}^*$  ▷ means and CCA projection matrices

4: if train_time then ▷ update statistics and CCA projections during training
5:    $\mu_x \leftarrow \frac{1}{m} \sum_i \mathbf{X}_i$  ▷ update  $\mu_x$  and  $\mu_y$  with means of batch
6:    $\mu_y \leftarrow \frac{1}{m} \sum_i \mathbf{Y}_i$ 
7:    $\bar{\mathbf{X}} = \mathbf{X} - \mu_x$  ▷ mean center data
8:    $\bar{\mathbf{Y}} = \mathbf{Y} - \mu_y$ 
9:    $\hat{\Sigma}_{xx} = \frac{1}{m-1} \bar{\mathbf{X}}' \bar{\mathbf{X}} + r\mathbf{I}$  ▷ estimate covariances of batch
10:   $\hat{\Sigma}_{yy} = \frac{1}{m-1} \bar{\mathbf{Y}}' \bar{\mathbf{Y}} + r\mathbf{I}$ 
11:   $\hat{\Sigma}_{xy} = \frac{1}{m-1} \bar{\mathbf{X}}' \bar{\mathbf{Y}}$ 
12:   $\mathbf{C}_{xx}^{-1} = \text{cholesky}(\hat{\Sigma}_{xx})^{-1}$  ▷ compute inverses of Cholesky factorizations
13:   $\mathbf{C}_{yy}^{-1} = \text{cholesky}(\hat{\Sigma}_{yy})^{-1}$ 
14:   $\mathbf{T} = \mathbf{C}_{xx}^{-1} \hat{\Sigma}_{xy} (\mathbf{C}_{yy}^{-1})'$  ▷ compute matrix  $\mathbf{T}$ 
15:   $\mathbf{e}, \mathbf{U} = \text{eigen}(\mathbf{T}\mathbf{T}')$  ▷ compute eigenvectors of  $\mathbf{T}\mathbf{T}'$  and  $\mathbf{T}'\mathbf{T}$ 
16:   $\mathbf{e}, \mathbf{V} = \text{eigen}(\mathbf{T}'\mathbf{T})$ 
17:   $\mathbf{A}^* \leftarrow \mathbf{C}_{xx}^{-1} \mathbf{U}$  ▷ compute and update CCA projection matrices
18:   $\mathbf{B}^* \leftarrow \mathbf{C}_{yy}^{-1} \mathbf{V}$ 
19:   $\mathbf{A}^* \leftarrow \mathbf{A}^* \cdot \text{sgn}(\text{diag}(\mathbf{A}^{*'} \hat{\Sigma}_{xy} \mathbf{B}^*))$  ▷ flip signs of projection matrices
20: else ▷ at test time use statistics estimated during training
21:    $\bar{\mathbf{X}} = \mathbf{X} - \mu_x$  ▷ mean center test data
22:    $\bar{\mathbf{Y}} = \mathbf{Y} - \mu_y$ 
23: end if
24:  $\mathbf{X}^* = \bar{\mathbf{X}} \mathbf{A}^*$  ▷ project latent representations with CCA projections
25:  $\mathbf{Y}^* = \bar{\mathbf{Y}} \mathbf{B}^*$ 

return  $\mathbf{X}^* \mathbf{Y}^*$ 
```

---

see that this does not necessarily lead to the best retrieval performance. The freely learned embedding  $Learned-\mathcal{L}_{rank}$  shows overall the lowest correlation but achieves comparable results to DCCA on this dataset. In terms of overall correlation,  $CCAL-\mathcal{L}_{rank}$  is situated in-between the two other approaches. We have seen in all our experiments that combining both concepts in a unified retrieval paradigm yields best retrieval performance over different application domains as well as data regimes. We take this as evidence that componentwise-correlated projections support cosine distance based embedding space learning.

## Architecture and Optimization

In the following we provide additional details for our experiments carried out in Section 5.

### Image-Text Retrieval

We start training with an initial learning rate of either 0.001 (all models on IAPR TC-12 and Flickr30k  $Learned-\mathcal{L}_{rank}$ ) or 0.002 (Flickr30k DCCA and  $CCAL-\mathcal{L}_{rank}$ )<sup>5</sup>. In addition, we apply 0.0001 L2 weight decay and set the batch size to 1000 for all models. The parameter  $\alpha$  of the ranking loss in Equation (8) is set to 0.5. After no improvement on the validation set for 50 epochs we divide the learning rate by 10 and reduce the patience to 10. This learning rate reduction is repeated three times.

### Audio-Sheet-Music Retrieval

Table 7 provides details on our audio-sheet-music retrieval architecture.

Table 7: Architecture of audio-sheet-music model. BN: Batch Normalization, ELU: Exponential Linear Unit, MP: Max Pooling, Conv(3, pad-1)-16:  $3 \times 3$  convolution, 16 feature maps and padding 1.

Sheet-Image $40 \times 100$	Spectrogram $136 \times 100$
$2 \times \text{Conv}(3, \text{pad-1})-16$ BN-ELU + MP(2)	$2 \times \text{Conv}(3, \text{pad-1})-16$ BN-ELU + MP(2)
$2 \times \text{Conv}(3, \text{pad-1})-32$ BN-ELU + MP(2)	$2 \times \text{Conv}(3, \text{pad-1})-32$ BN-ELU + MP(2)
$2 \times \text{Conv}(3, \text{pad-1})-64$ BN-ELU + MP(2)	$2 \times \text{Conv}(3, \text{pad-1})-64$ BN-ELU + MP(2)
$2 \times \text{Conv}(3, \text{pad-1})-64$ BN-ELU + MP(2)	$2 \times \text{Conv}(3, \text{pad-1})-64$ BN-ELU + MP(2)
Conv(1, pad-0)-32-BN	Conv(1, pad-0)-32-BN
GlobalAveragePooling	GlobalAveragePooling
Respective Optimization Target	

As in the experiments on images and text we optimize our networks using *adam* with an initial learning rate of 0.001 and batch size 1000. The refinement strategy is the same but no weight decay is applied

<sup>5</sup>The initial learning rate and parameter  $\alpha$  are determined by grid search on the evaluation measure MRR on the validation set.

and the margin parameter  $\alpha$  of the ranking loss is set to 0.7.

### Zero-Shot Retrieval

Table 8 and 9 provide details on the architectures used for our zero-shot retrieval experiments carried out in Section 5.4. The general architectures follow Reed et al. [25] but are optimized with a pairwise ranking loss in combination with our proposed CCA layer. The dimensionality of the retrieval space is fixed to 64 and both models are again optimized with *adam* and a batch size of 1000. The learning rate is set to 0.0007 for the CNN and 0.01 for the CRNN and. The margin parameter  $\alpha$  of the ranking loss is set to 0.2. In addition we apply a weight decay of 0.0001 on all trainable parameters of the network for regularization.

Table 8: Architecture of Zero-Shot Retrieval CNN. VS: Vocabulary Size, BN: Batch Normalization, ELU: Exponential Linear Unit, MP: Max Pooling, Conv(3, pad-1)-16:  $3 \times 3$  convolution, 16 feature maps and padding 1.

ImagenNet Feature 1024	Text $VS \times 30 \times 1$
FC(1024)-BN-ELU	$1 \times \text{Conv}(3, \text{pad-}same)-256$
FC(1024)-BN-ELU	BN-ELU + MP(3, 1)
FC(64)	$2 \times \text{Conv}(3, \text{pad-}valid)-256$
	FC(1024)-BN-ELU
	FC(64)
Respective Optimization Target	

Table 9: Architecture of Zero-Shot Retrieval CRNN. VS: Vocabulary Size, BN: Batch Normalization, ELU: Exponential Linear Unit, MP: Max Pooling, Conv(3, pad-1)-16:  $3 \times 3$  convolution, 16 feature maps and padding 1. GRU-RNN: Gated Recurrent Unit [5]

ImagenNet Feature 1024	Text $VS \times 30 \times 1$
FC(1024)-BN-ELU	$1 \times \text{Conv}(3, \text{pad-}same)-256$
FC(1024)-BN-ELU	BN-ELU + MP(3, 1)
FC(64)	$2 \times \text{Conv}(3, \text{pad-}valid)-256$
	GRU-RNN(512)
	TemporalAveragePooling
	FC(64)
Respective Optimization Target	